Collaborative Medical Informatics Research Using the Internet and the World Wide Web

Edward H. Shortliffe, MD, PhD[1], G. Octo Barnett, MD[2], James J. Cimino, MD[3], Robert A. Greenes, MD, PhD[4], Stanley M. Huff, MD[5], and Vimla L. Patel, PhD[6]

[1]Section on Medical Informatics, Stanford University, Stanford, California
[2]Laboratory of Computer Science, Massachusetts General Hospital, Boston, Massachusetts
[3]Department of Medical Informatics, Columbia University, New York, New York
[4]Decision Systems Group, Brigham and Women's Hospital, Boston, Massachusetts
[5]Department of Medical Informatics, University of Utah, Salt Lake City, Utah
[6]McGill Cognitive Science Centre, McGill University, Montreal, Quebec

The InterMed Collaboratory is an interdisciplinary project involving six participating medical institutions. There are two broad mandates for the effort. The first is to further the development, sharing, and demonstration of numerous software and system components, data sets, procedures and tools that will facilitate the collaborations and support the application goals of these projects. The second is to provide a distributed suite of clinical applications, guidelines, and knowledge-bases for clinical, educational, and administrative purposes. To define the interactions among the components, datasets, procedures, and tools that we are producing and sharing, we have identified a model composed of seven tiers, each of which supports the levels above it. In this paper we briefly describe those tiers and the nature of the collaborative process with which we have experimented.

INTRODUCTION

The idea of an electronically mediated research collaboration, a "collaboratory," was succinctly defined by Dr. William Wulf [1]: "A national collaboratory is a center without walls, in which the nation's researchers can perform their research without regard to geographical location—interacting with colleagues, accessing instrumentation, sharing data and computational resources, [and] accessing information in digital libraries... [Such a collaboratory is] more than a mere interconnection of computers [and offers] a complete infrastructure of software, hardware, and networked resources to enable a full range of collaborative work among scientists." Taking advantage of the rapid change in computing and communications technologies, the InterMed project seeks to demonstrate the viability of the collaboratory concept in the context of medical informatics research. InterMed was initiated as an NLM-funded collaboration among Stanford's Section on Medical Informatics, the Harvard Decision Systems Group (DSG), and the Columbia Presbyterian Medical Center (CPMC). Workers at Massachusetts General Hospital, the University of Utah, and McGill University are also closely involved. The broad project goals are:

- to accelerate our individual progress in building advanced clinical, educational, and research applications using software and information components made available by our collaborators;
- to build new biomedical and clinical applications using shared components and methodologies;
- to provide a broadly applicable model for such collaborative work;
- to evaluate the collaboratory activities with a variety of functional, cognitive, and observational metrics.

We envision these efforts as precursors to a national (or international) collaboratory with many more participants. Thus scalability is a prime concern.

MOTIVATION

Collaboration is often difficult, even within institutions where proximity may not overcome problems with "chemistry" among individuals, competing methodologies, or personality clashes. It is accordingly appealing to seek collaborations with individuals of like interests and commitment, even if there is a substantial geographic barrier to be overcome. The Internet's capabilities offer an exciting mechanism for making this work. Our six institutions, and our past research experiences, are complementary rather than duplicative. Despite a previous tradition of separateness, we have been encouraged by recent progress because joint projects aimed at component development and integration are proceeding with the Internet as catalyst. This paper provides a status report on the collaboration itself; reports on individual subprojects are cited below.

COLLABORATIVE ARCHITECTURE

The model on which we have built our collaboration involves the decomposition of current development activities into sets of components that can be constructed at one or more institutions, shared with

0195—4210/96/$5.00 © 1996 AMIA, Inc.

others, and then melded to provide robust applications that could not be developed as effectively by any single group. We have developed a model composed of 7 tiers (Figure 1), each of which supports the levels above it: (1) Infrastructure and Services: The network itself, plus the associated systems methodologies; (2) Vocabulary/Taxonomy: The vocabularies needed to support components and applications; (3) Data and Knowledge: The databases and knowledge bases on which clinical applications are built; (4) Agents and Components: The elements that build on tiers 1 through 3 and support modular construction of clinical systems; (5) Composition/Development Environments: The application frameworks and development tools that depend on the components and agents while providing for creation of applications; (6) Testbed Applications: Applications to support clinical care, education, and decision making; and (7) Collaborative Policies: The policies, principles, and guidelines (regarding ownership, liability, intellectual property, confidentiality, security, and the like) that must be developed and implemented in defining models by which applications are constructed, and software is shared, across communication networks.

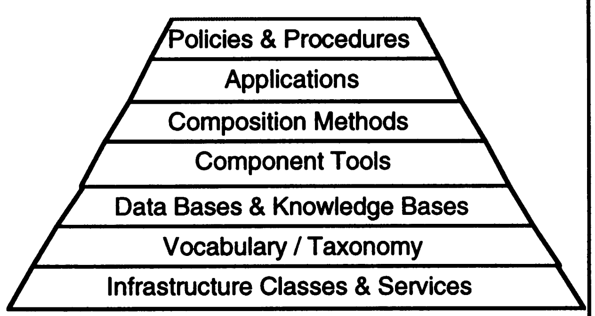

Figure 1. 7-Tiered Collaborative Architecture

DOMAIN-SPECIFIC MIDDLEWARE

The seven-tiered collaborative model is perhaps best understood in the context of evolving notions of *middleware* that have been extensively discussed in the computer-science and telecommunications communities [2].

For the Internet today, there is an underlying, readily available, infrastructure which consists of elements such as:

- Networks, workstations, servers
- Internet connectivity
- Network-aware operating systems with network stack software
- Protocol suites (TELNET, FTP, HTTP, SMTP, IMAP, NNTP, Z39.50, Gopher, multicast, other special protocols as needed)
- Environment-specific software tools
- Information services (interest lists, news groups, FTP archives, gopher archives)

On top of this infrastructure, people build various stand-alone applications—some general and some for particular domains—all essentially from scratch. There is very little code shared among applications and they communicate with each other by means of copy/paste or file interfaces in most circumstances. These vertically-oriented application systems require a great deal of custom-tailored development effort.

We (and many others who are concerned about the complexity of software development, the time required to field new applications, and the failure to benefit from the experiences of others except through published articles) support ongoing efforts to develop additional *middleware* layers that need to be added to the network's architecture (Figure 2). There are at least two kinds of such middleware under development: *general* middleware (sitting atop the existing infrastructure and providing new kinds of services across application domains) and *domain-specific* middleware. General middleware, for example, includes general distributed-computing tools that interoperate among heterogeneous platforms—elements such as:

- Software development tools, libraries, prototyping tools
- Multimedia tools
- Federated database tools
- Collaboration tools and groupware
- Reusable modules that implement protocol client or server functions for incorporation in agents or other applications (e.g., email responder tools, Gopher protocol handler, KIF/KQML handlers)
- Data and knowledge interchange tools
- Object Management Group (OMG) tools and reusable object libraries

In addition to this general layer, there needs to be a layer of tools and services that cater to the domain-specific needs of biomedical applications (*biomedical* middleware) and that facilitate the development of more horizontally integrated workstation environments for medical professionals. This layer may include elements such as:

- Special object libraries for clinical data access, image manipulation,
- Tools for creating and accessing domain data structures (e.g., vocabularies, patient record databases, clinical-trials methods, protocols and practice guidelines, images)

- Tools for supporting domain services (e.g., protocol eligibility, protocol treatment, progress-note-entry forms content)

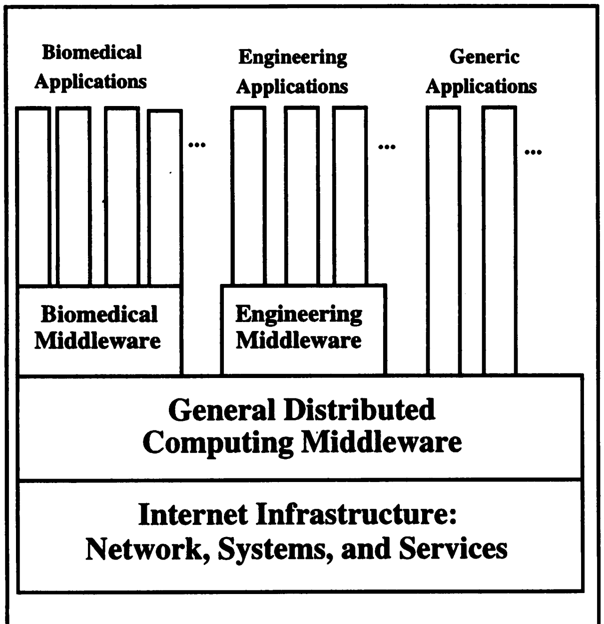

Figure 2. Generic and Domain-Specific Middleware

With such middleware in place, applications should be easier to develop and have a more horizontal integration that supports all the particular needs of various classes of biomedical users. As experimental elements demonstrate utility and gain acceptance, they may move out of the applications and into middleware below. Thus one goal of our work is to focus on notions of biomedical middleware development while recognizing that specific applications from above will drive the software design and the functionalities that we seek..

VOCABULARY SERVER WORK

During the first two years of InterMed work, much of the emphasis has been placed on developing a shared view of how a generic clinical vocabulary should be structured (Tier 2) [3]. We began by investigating representation tools to support developing, browsing, and modifying vocabulary knowledge structures. We surveyed the InterMed sites to focus the collaborative vocabulary server work on specific application goals stimulated by ongoing projects at the various institutions. Using an available web-based knowledge-representation tool known as Ontolingua[4], we have adapted an Ontolingua WWW-based vocabulary editing tool for collaborative use among the InterMed project members [5, 6]. The editor has two immediate benefits: (1) because it is web-based, it is hardware independent and usable on any machine that can run standard web browsers, and (2) it shields developers from the syntactic details of the Ontolingua language, allowing editing of the vocabulary content without requiring knowledge of Ontolingua (or even a local installation of Ontolingua). A software tool automatically generates HTML pages from Ontolingua and is used to update WWW pages on the server at Stanford as workers at Columbia, for example, add to the content of the vocabulary.

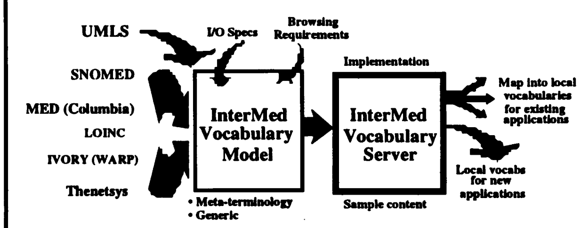

Figure 3. Development of a Collaborative Model for a Clinical Vocabulary

Based on our early work, and guided by intense study of each other's vocabulary conventions, by the summer of 1995 we had defined a process for developing a shared vocabulary model as is shown in Figure 3. We were influenced by the major terminology effort coordinated by the National Library of Medicine (the Unified Medical Language System, or UMLS, with which several InterMed collaborators had also been involved as developers). We also sought to understand overlapping and conflicting needs and features of clinical vocabularies that were in use at the various sites. These included SNOMED, which had been studied extensively and adopted in part by many of us, but also the Medical Entities Dictionary (MED) which is in use at Columbia (we created a version of Columbia's MED in Ontolingua on the Stanford server—known as "The InterMED"), the LOINC system (a national development effort with which the University of Utah has been closely involved), the IVORY vocabulary (which evolved from Wisconsin's WARP terminology and has been used for the T-HELPER AIDS management project at Stanford), and Thenetsys vocabulary in use at Brigham and Women's Hospital. Guided by an analysis of the input/output characteristics that would be required by a network-based server if it were to meet the needs of our various application projects, as well as the associated browsing/maintenance requirements, we sought to develop a generic, implementation-independent vocabulary model which used a set of "meta descriptors" to define semantic elements that need to be represented [3]. Using this model, which continues to evolve as we learn more about each other's work and requirements, an implementation has been created using the Ontolingua server mentioned above [7]. Our goal is to demonstrate that the server can in turn provide the basis for mapping sub-vocabularies, represented using the generic model, back into the local vocabularies that we require for specific applications at our individual sites. In this way, future vocabulary-development efforts, at our institutions and elsewhere, may be able to draw on the generic model and server as they are fleshed out with additional content. The effort is not meant to duplicate or compete with existing vocabulary efforts but, rather, to provide a generic model for how they can be related to one another, with the UMLS playing a key lookup role as the master thesaurus.

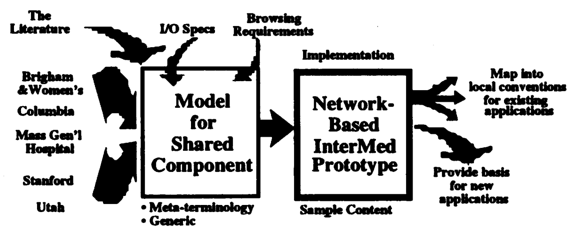

Figure 4. Generic Approach to Collaborative Model Development

COLLABORATIVE MODEL DEVELOPMENT

As we gained insight into the process by which we were evolving a shared view of the vocabulary server's role and representation requirements, we attempted to abstract the process, defining an approach to "collaborative model development" which we hoped would apply to other multi-institutional group development activities (Figure 4). More recently we have linked the vocabulary development work to the generation of a shared model for clinical guideline representation and use (Tiers 3 and 4), recognizing that a set of applications related to guidelines would be extremely helpful in testing/evaluating the relevant facilities of our vocabulary server while providing the focusing power of specific performance expectations. All InterMed sites are working locally on guidelines development and integration of guidelines into clinical systems (at Tier 6). Through extensive discussion and email exchanges, we have begun to identify the common characteristics as well as methods for resolving some of the apparent central differences among the approaches. We are applying our generic approach for collaborative model development to the notion of a network-based guidelines server [8].

MECHANISMS FOR COLLABORATION

One of the goals of the InterMed project has been to use advanced communications technologies to assess the strengths and limitations of network-mediated collaboration, distinguishing true collaboration (in which the results emerge from shared contributions from several participants, making it difficult to assign credit for the total product to any individual or individuals) from mere cooperation. Throughout the project, collaborators have been in almost daily contact, largely via electronic mail, and our online email archives are extensive. The archives also provide a detailed record of both the form and content of the joint design and development activities that have been assumed under the InterMed umbrella.

We learned early in our work that email's narrow bandwidth is especially effective for communication among individuals who know each other well and have shared perceptions of the problems on which they are jointly working. To try to address this issue, we have initiated occasional face-to-face meetings during national meetings or when members of one institution have visited groups at other institutions to see their systems in operation, to examine data structures in detail, and to get a much stronger understanding of how the developing shared resources need to be configured in order to meet the broad requirements of the InterMed community. After these visits, the email communication among the groups has become more concrete and effective.

In addition, teleconferencing (rather than video-conferencing, which proved expensive and unnecessary) has been enthusiastically accepted. We hold conference calls for 90 minutes every two weeks, involving essentially all project members at all six sites. The addition of McGill cognitive scientists has added an new dimension to the collaborative work. The McGill team largely listens (and records for later transcription) during the conference calls, viewing the discussions as experimental material for analysis as we attempt to measure and evaluate the collaborative process (see below).

An additional technology warrants mention. In preparation for a 1995 InterMed demonstration session[7], we needed to coordinate contributions from each of the collaborative sites. This required extensive email exchanges but also "distributed group meetings" at which we were able to practice talks, show planned slides, and demonstrate software to one another. To support demonstrations, we successfully experimented with Timbuktu, a cross-platform software tool that allowed machines at each site, hooked up to projection devices, to be linked over the Internet. One machine would "drive" the displays on all linked machines, thereby allowing each site to run its own demo while the researchers at the other sites observed and commented. This software provides useful additional support for the kinds of collaborative processes with which we are experimenting.

EVALUATION PROCESS

The Internet is well established as a tool for collaborative research in areas such as nuclear physics and biotechnology. With its increasing role in medicine, it has been natural for the Internet to become a tool for collaborative medical informatics research. We seek to understand which aspects are successful and which fail. We accordingly are evaluating the InterMed Collaboratory with two complementary objectives in mind. The first is to undertake a formal formative evaluation process, in which we are assessing InterMed in terms of how successful the collaborative project has been in attaining its stated objectives. The second evaluation focuses on more fundamental research issues. The objective is to evaluate InterMed as an ongoing experiment in computer-mediated collaborative design, and to understand what we can learn about this process and about the Internet as a medium for such an endeavor. The objectives are interdependent, and the research-oriented evaluation process is providing a basis for developing and continually refining guidelines for formal evaluation. The analytic framework is also facilitating the ways in which we can contribute to the iterative design-implementation-evaluation process. Early results of this evaluation work have been reported elsewhere [9].

CONCLUSIONS

Among the InterMed lessons are: (1) the InterMed 7-tiered architecture has proved to be a useful model for structuring our collaborative development activities and the sharing of components and tools; (2) the Web is an effective mechanism for sharing rapidly changing software, data structures, reports, and diagrams among geographically distributed workers who use a variety of computing platforms; (3) email is an effective means of collaborative communication only after the groundwork has been laid with detailed discussions and development of a shared culture among the participants; and (4) current video conferencing is not effective given the limited quality of the images and the expense associated with its use. We are using audio conferencing for group meetings on a regular basis, augmenting the phone calls with Timbuktu interactions when shared demonstrations of systems are important to the purposes of the call. The combination of email, audio conferences, occasional meetings, and shared interests plus a commitment to collaboration have combined to facilitate the goals of the InterMed collaboratory.

ACKNOWLEDGMENTS

We gratefully acknowledge the contributions and support of InterMed team members at all six participating institutions. This work has been supported in part by High Performance Computing and Communications contracts from the National Library of Medicine to Brigham and Women's Hospital, Columbia University, and Stanford University, with subcontracts to Massachusetts General Hospital and the University of Utah. Further information regarding the InterMed Collaboratory may be found at the following URLs:

Brigham and Women's Hospital / Harvard
http://dsg.harvard.edu/public/intermed/InterMed_Collab.html

Columbia Presbyterian Medical Center
http://www.cpmc.columbia.edu/intermed_proj.html

Massachusetts General Hospital
http://lcs-guide.mgh.harvard.edu/

Stanford University School of Medicine
http://www-smi.stanford.edu/projects/intermed-web

REFERENCES

1. Committee on a National Collaboratory, Computer Science and Telecommunications Board, National Research Council. *National Collaboratories: Applying Information Technology for Scientific Research*. Washington, D.C.: National Academy Press, 1993.

2. Korzeniowski P. Make way for data. Byte 1993;18(7):113-116.

3. Oliver D. *Collaborative Development of the InterMed Vocabulary Model.* Section on Medical Informatics, Stanford University. Internal Technical Report August 1995.

4. Farquhar A, Fikes R, Pratt W, Rice J. *Collaborative ontology construction for information integration.* Knowledge Systems Laboratory, Computer Science Department, Stanford University. Technical Report KSL-95-63. August 1995.

5. Gennari J, Oliver D, Pratt W, Rice J, Musen M. A Web-based architecture for a medical vocabulary server. *Proceedings of the Nineteenth Annual Symposium on Computer Applications in Medical Care*. New Orleans, LA: American Medical Informatics Association, 1995:275-279.

6. Rice J, Farquhar A, Piernot P, Gruber T. Using the Web instead of a window system. *Proceedings of Conference on Human Factors in Computer Systems (CHI '96)*. Vancouver, B.C., Canada, 1996.

7. Oliver D, Barnes M, Barnett G, et al. InterMed: An Internet-based medical collaboratory [demo abstract]. *Proceedings of the Nineteenth Annual Symposium on Computer Applications in Medical Care*. New Orleans, LA: American Medical Informatics Association, 1995:1023.

8. Pattison-Gordon E, Cimino J, Hripcsak G, et al. Requirements of a sharable guideline representation for computer applications. Technical Report. *Decision Systems Laboratory,* Brigham&Women's Hospital, Boston, MA., April 1996.

9. Patel V, Cimino J, Shortliffe E. Explorations in Internet-based collaborative informatics research—A cognitive evaluation [abstract]. In: Humphreys B, ed. *Spring Congress of the American Medical Informatics Association*. Kansas City, MO: American Medical Informatics Association, 1996.